arXiv:1703.07075v1 [cs.AI] 21 Mar 2017

Pseudorehearsal in value function approximation

Vladimir Marochko, Leonard Johard, and Manuel Mazzara

Innopolis University
Universitetskaya str. 1, 420500 Innopolis, Russia

Abstract. Catastrophic forgetting is of special importance in reinforcement learning, as the data distribution is generally non-stationary over time. We study and compare several pseudorehearsal approaches for Q-learning with function approximation in a pole balancing task. We have found that pseudorehearsal seems to assist learning even in such very simple problems, given proper initialization of the rehearsal parameters.

Keywords: Reinforcement learning, rehearsal, pseudorehearsal, catastrophic forgetting

1 Introduction

Reinforcement learning is a more general problem formulation than the commonly used supervised learning framework. As such, it can be applied to a wider range of problems. It is also a more difficult problem to optimize for, as the feedback is more limited.

Reinforcement learning covers space of prediction and control problems in partially unknown space with unknown behavior. The agent should explore the environment and find optimal actions it has to perform to reach the goal. It is used in cases when the optimal policy is unknown so there is no way to train agents using supervised learning algorithms. The basic idea can be described by the metaphor of a player starting an unknown game and, after a number of turns, he/she receives a message stating that he/she has lost or won. After a number of games he/she will figure out how to act to win as often as possible.

1.1 Supervised agents in reinforcement learning problems

In order to solve the practical problems that can be assumed to be approximately Markov decision process (MDP) [1], like robot's navigation, playing chess or trading on stock exchange, we can use a value function approximation to speed up the learning process [2] [3]. The weakness of this approach is that convergence is not guaranteed if the MDP approximation is incorrect, or in cases where the inputs are continuous, and which necessitates non-linear function approximation [4]. Furthermore if we have continuous outputs the value approximation needs to be combined with an additional optimization technique, such as REINFORCE [5], in order to search for optimal outputs. If the outputs are discrete, a simple

maximization is sufficient and allows us to make use of the Q-learning framework [6][7].

Although a value function simplifies the learning problem by effectively converting a reinforcement learning problem to a supervised learning problem through the use of bootstrapping [8], it is still a more difficult problem than conventional supervised learning. One of these additional difficulties is that the policy-dependent rewards introduces a concept drift [9]. This introduces the risk of unstable oscillations, but is generally solvable at the cost of slower learning if the learning rate is set sufficiently small in the beginning. This ill conditioning of the problem has been considered one of the main challenges of reinforcement learning.

1.2 Catastrophic forgetting

A different problem which recently got more attention is catastrophic forgetting [10]. This problem is most commonly described in an online unsupervised Hebbian learning task, where the ability to retrieve previously stored patterns is lost as we update weights in the training of new patterns.This catastrophic forgetting of the original patterns takes place even when parameter space is more than sufficient to store both sets of patterns and is a consequence of the limited mixing of input objects.

Simply alternating between the new and old patterns group with a sufficiently low learning rate would in theory solve the problem, but this has a potential impact on the convergence rate and requires explicit memorization of all training patterns. In order to minimize the effect on convergence rate, we would like to maximize the mixing of the presented inputs.

2 Catastrophic forgetting in reinforcement learning

The online nature of reinforcement learning means that catastrophic forgetting is a key bottleneck. There are two principal non-sharpening approaches to the catastrophic forgetting problem suggested in literature: rehearsal and pseudorehearsal. In addition, other methods based on sparse representations [11] [12] have been used less frequently. This latter approach has a theoretical downside in its negative impact on the ability to generalize, but has shown at least mixed results and is possible to use in conjunction with the rehearsal methods.

2.1 Rehearsal approaches

The first and most straight-forward principal approach for mitigating catastrophic forgetting is rehearsal [13], [14]. A rehearsal strategy simply stores a subset of all previous experiences in a buffer. When a new pattern is presented, this pattern is combined with several patterns from the buffer in order to form a learning batch with good mixing. There are several possible heuristics for selecting patterns for rehearsal.

The importance of catastrophic forgetting in reinforcement learning was identified early. Lin introduced the term Experience Replay [15] for referring to the

use of rehearsal strategies in the reinforcement learning setting. Such rehearsal has shown very promising results in robotics [16] and on more complex environments, such as Deep Q-learning for playing Atari games [17].

2.2 Pseudorehearsal approaches

A second principal approach to solving catastrophic forgetting is pseudorehearsal [18], which does not require explicit storage of patterns. Instead, it uses a two-step process where generative models are learnt alongside with the main task. These generative models create pseudopatterns, which are combined in batches with real patterns for training the agent.

An interesting questions is whether these generated approximations of the real data are sufficiently accurate in practice to reduce forgetting. Remarkably, even extremely crude generative models have proven highly effective. In the original work in this area by [18], pure noise fed to the network was able to almost completely eliminate catastrophic interference. The argument of the authors was that, although the input is completely random, the activation distributions in deeper levels of the network will be representative of the learnt input data.

An analytical approach by Frean and Robins [19] in single perceptrons suggest an alternative explanation for the surprising efficiency of random pseudopatterns. They suggest that the pseudopatterns approximate the mean of the input. Training on this mean of the input leads to decorrelation of the input patterns, which in high dimensional inputs makes the different patterns' weight updates orthogonal to each other. In addition, they demonstrated that using this mean directly was at least efficient as generating pseudopatterns. Further work in this direction was done in a thesis by Goodrich [20], where some of these results where expanded to multilayer perceptrons.

Regardless of the reason for such networks, pseudorehearsal methods have been demonstrated to significantly decrease and almost completely eliminate the catastrophic forgetting in unsupervised learning [18], supervised learning [13] and reinforcement learning [21]. It is interesting to note that the results of Baddeley suggest that the widely studied ill conditioning might not be the main bottleneck of reinforcement learning after all. Instead, their results indicate that the catastrophic forgetting is the main bottleneck for reinforcement learning problems.

Pseudorehearsal algorithms For testing pseudorehearsal approach we used two different pseudorehearsal types and the online learning with one backpropagation step as an example of learning without pseudorehearsal. One algorithm is based on correcting of the weight updates, other is a batch-backpropagation learning.

The first pseudorehearsal algorithm is the one used by Frean and Robins [19] with a simplified weighting equation and changed for non-linear neural network inner assignments. The idea behind algorithm is to generate pseudoset, feed it through the network and save activations on each neuron for every pseudoexample. Then the agent is learned online, but when the real example fed to the network we use equation

$$\Delta w_i = err_{b_i} \frac{1}{pr} \sum_{j=1}^{pr} \frac{b_i x_{ij} \cdot x_{ij} - x_{ij} x_{ij} \cdot b_i}{b_i \cdot b_i x_{ij} \cdot x_{ij} - b_i \cdot x_{ij} b_i \cdot x_{ij}}$$

to update w_i - weights at the i^{th} layer, where err_{b_i} - vector backpropagation errors of the learned example at the i^{th} layer, pr - size of pseudoset, b_i - vector of activations of the learned example when fed forward through the network and x_{ij} is vector of activations of the j^{th} pseudoset on the i^{th} layer.

The second one - is straight-forward using pseudosets in batch backpropagation learning - we generate set of pseudoexamples, feed it through the network, save the network outputs as the targets and then create a matrix of feature vectors where first vector is real example, others are pseudovectors, and matrix of targets where first vector is target for the real example and others are saved earlier networks outputs on pseudoset, then each time we learn agent on the whole set.

2.3 Biological forgetting

An interesting particular case of catastrophic forgetting problem is learning in the human brain. Dual network models were initially inspired by biological learning [22]. As a consequence of promising experimental results of such networks, pseudorehearsal was indeed found to be the most plausible explanation for the otherwise cryptic need for dual learning systems in the brain [23].

More biologically detailed extensions of these models have recently been explored by Hattori, where they again showed excellent improvements on the ability to store information [24].

The pseudorehearsal approach also contributes to the urgent need for new biological plasticity rules in large scale neurosimulation and especially for their developmental varieties (e.g. BioDynaMo [25]). Real full scale brain simulation is approaching, but we are still lacking even a basic understanding of the role dreams play in the learning process. This despite the fact that sleep stages are of considerable length and evident in even the simplest of biological neural networks.

3 Experimental design

We will reevaluate the analytic results of Frean and Robins [19] in a real reinforcement learning task. We evaluate and compare two algorithms for pseudorehearsal on a pole balancing task using Q-learning with function approximation. Further, we will study the effect of sparsity in fulfilling the requirement for a high dimensional input space these algorithms relied on. Our agent is a classic Q-learning agent with ϵ-greedy policy using a feed-forward-backpropagation neural network as function approximator, discounted factor of agent is 0.9. The environments used for training is the single-pole balancing cart.

Observation Two different observations are used for experimental comparison. The first observation type given to the agents constitute a fully observable MDP

and includes current position, velocity, acceleration of the cart, as well as the current angle, angular velocity and angular acceleration of the pole. The second observation type make the problem partially observable - here the agent knows only cart's position and pole's angle. If the cart reaches the end of track or the pole falls for angle more than predefined pole failing angle - the game is lost and the agent is gained negative reward for this task.

We represented the observations as a feature vector by two different methods. The first method was a representation of input values where the i-th observation value sets the $2 * i$-th feature if it is positive or on $2 * i + 1$-th if it is negative. The second method was to use sparse unary vectors where feature vector is concatenation of parameter vectors, similar to a table. Each of parameter vectors consisted of elements associated with discrete values inside the range possible for each parameter - $[-20; 20]$ for linear parameters, $[-60; 60]$ for angular. All the elements of the vector were set to zeros except two - the element associated with the rounded value of the parameter was set to one, and the next element is set to the fractional part of this parameter.

Performance metric Agent tries to balance pole or poles as long as it can for 5000 tries, for two sets of parameters we also made an averaged variants where 10 iterations of this 5000 tries are averaged to make sure that convergence tendency is reproducible and not a set of random successful moves. We also have results for fully random policy to compare with.

Parameter settings The task was repeated with different sizes of pseudoitem batches, with different numbers of iterations between reinitialization of the pseudosets and with different learning rates. The learning rates used were 0.1, 0.01 and 0.001. The discount factor was set to 0.9. The sizes of pseudosets were 10, 30, 50 and 100 pseudoitems, respectively. We resample a new set of pseudoitems after every 1, 10 or 100 runs. Parameters are chosen to define influence of size of the pseudoset and frequency of it's reinitialisation on learning. We try to cover a wide range without trying all the possible values. And we decrease parameters in close to geometrical progression to see if the influence is logarithmic. For 30 and 50 item pseudorehearsal batches we also tried 30 and 50 reinitialisation gaps to increase coverage.

Performance metrics We did not stop learning after an agent reaches satisfactory result, because the continued learning contains cases of catastrophic forgetting which we would like to explore. As we had a very short learning time during which the performance increases followed by long row of tries with unstable behavior, we evaluate the efficiency of different approaches by measuring mean and median number of steps per try for each approach and compare these two numbers.

$Mean > median$ indicates that some agent's tries were highly effective and agent balanced pole for a long time, while the most of runs were weak, so no convergence occurred or the influence of catastrophic interference is too high to handle needed weights for a long time. $Mean < median$ shows that agent successfully converge to some optimal policy, and it's policy is stable but some tries are failed so bad that it affected the whole picture, so in this case we can

see successful learning, strong influence of catastrophic forgetting and successful avoidance of this influence. $Mean \approx meadian$ means that catastrophic forgetting and its avoidance has nearly the same influence. Results were averaged over ten

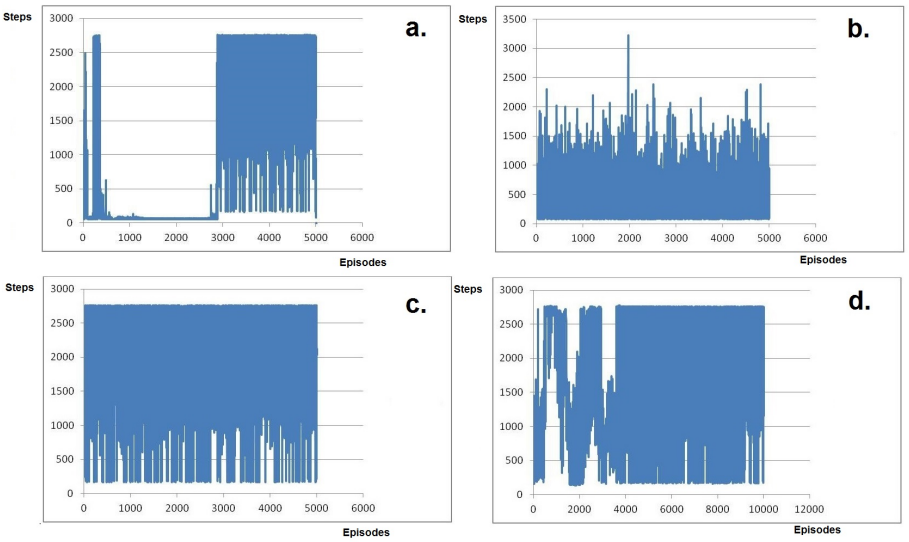

Fig. 1. Examples: a. median $\ll$ mean; b. median $<$ mean; c. median $>$mean; d. median $\approx$ mean

runs when training times allowed and other cases we presents results over single runs.

3.1 Results

Results for all observations show different learning for different cases, while averaged approach and comparison with the random agent assure us that the learning has place and in case without pseudorehearsal it depends on learning rate mostly. For the agents using pseudorehearsal learning depends on sets of parameters - learning rate, pseudoset size and relearning gap. Some of them can make agent to balance pole for significantly larger time than a random run, while the others perform the same, or worse, or even worse than a random agent.

For both used metrics we discovered a roughly bell-shaped graph of dependencies for each of used learning rates and for each of used techniques. All pseudorehearsal approaches have different sets of parameters providing optimal learning, a suboptimal and worse - in case of parameters a little different from the optimal ones. Further away from optimal parameters agents started to diverge, e.g. tried to drop the pole about four times faster than if it would use random policy, etc.

Both Frean-Robins and batch approaches perform similarly in for this observation, but their optimal parameters differ for the same learning rates. All this results are summarized in table 1 for MDPs and table 2 for partially observable MDPs - POMDPs.

3.2 MDP

For the cases of fully observable MDP, where agent knows anything about the current state of the pole cart and the pole the learning time when the agent's performance goes from some initial random to the final one is very short, agent quickly converges at some number of steps it can balance and most of its next moves holds around this result with some deviations, sometimes very large, caused by catastrophic forgetting. If some learning case makes agent signifi-

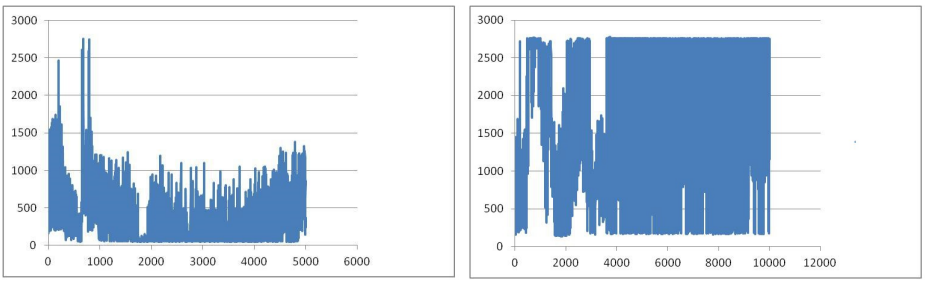

Fig. 2. left plot - learning rate = 0.001, batch learning, pseudoset size =10, relearning gap = 10; right - learning rate = 0.01, pseudoset size = 100, relearning gap = 100

cantly change its behaviour - change is as quick as initial learning and on graph looks like immediate change of performance.

3.3 POMDP

Partial observability tends to suffer less from forgetting, possibly because each part of the smaller space are more frequently visited. Pseudorehearsal has a more significant impact here: although it shares the same optimal-suboptimal-worst sets of parameters, optimal ones further decrease the number of runs needed to learn and decrease influence of the catastrophic forgetting: if the agent in current set of parameters doesn't diverges to the worst possible case - it's median for all runs is always higher then the mean, indicating a relatively stable behaviour after training. On the other hand, agents in POMDPs agents that can't converge change their policy more frequently than agents in MDPs, and while the agent in fully observable MDP has a minor chance to reach good performance after it reached a suboptimal solution, agent in POMDP easily changes its policy both ways.

Random		mean	217.72							
		median	206							
Learning rate	Fully observable 0.1					Fully observable 0.01				
No PR1	767.7612					525.7532				
	460					398				
No PR2	653.255					485.9627				
	355					370				
Frean-Robins	pseudoset	10	30	50	100	pseudoset	10	30	50	100
	relearn					relearn				
	1	88.4	479			1	53.8746	1425.314		
		84	339				53	1209		
	10	198.9974	260	1784	288.7504	10	81.161	61.0631	955.4878	107.554
		119	211	1938.5	210		54	54	1060	81
	30		573			30		460.5927		
			442					415		
	50			655		50			168.7599	
				276.5					66	
	100	722	413		1982.6	100	95.4113	272.1958		178.9365
		297	217		2058		82	107		113
batch	pseudoset	10	30	50	100	pseudoset	10	30	50	100
	relearn					relearn				
	1	181.9208	234.434			1	66.2465	388.2868		
		53	68				53	244.5		
	10	578.7582	159.062	562.722	92.5414	10	113.2239	73.4251	525.2255	1853.865
		119	89	324.5	55		53	53	400	1873
	30		1085.417			30		64.727		
			725					53		
	50			503.2082		50			2032.241	
				91					2142.5	
	100	389.2236	747.8532		56.3434	100	85.1528	178.7005		1065.519
		155	383		53		56	92		651
	Fully observable 0.001					Fully observable 0.1 averaged				
	316.564					1561.47782				
	169					1569.55				
	528.7564					1561.47782				
	398					1569.55				
Frean-Robins	pseudoset	10	30	50	100	pseudoset	10	30	50	100
	relearn					relearn				
	1	117.7718	251.8404			1	608.9561	433.9956		
		92	163				609.25	423.5		
	10	84.1524	110.2418	73.33	694.8546	10	460.5518	859.295	705.3484	309.2548
		56	86	79	663.5		467.15	862	696.7	296.85
	30		549.4772			30		447.0809	284.7513	
			434					439.7	262.4	
	50			54.258	309.9588	50				
				53	76					
	100	513.2786	74.9174			100	423.9774	677.3001		481.9075
		383	75				413.65	679.85		480.55
batch	pseudoset	10	30	50	100	pseudoset	10	30	50	100
	relearn					relearn				
	1	409.2142	1114.044			1	447.7925			
		299	818				420.1			
	10	369.9044	290.125	76.6654	253.8526	10	500.5673	542.4561		471.3174
		238	84	56	103		487.55	516.05		454.15
	30		410.3088			30		492.2506		
			325					465.35		
	50			267.8876		50				
				191						
	100	72.6846	177.4432		278.7542	100	381.1586			481.4641
		54	73		180		363.9			463.9

Table 1. MDP results overall

Learning rate	Fully observable, sparse vector, 0.01					Partially observable, 0.01					Partially observable, 0.01, averaged by 10 runs				
No PR1 mean	233.8168					1856.8124					617				
median	54					1961					610				
No PR2 mean	271.5718					1422.7302					880				
median	54					1348.5					880				
Frean-Robins	pseudoset	10	30	50	100	pseudoset	10	30	50	100	pseudoset	10	30	50	100
	relearn					relearn					relearn				
	1	178.994	182.0386			1	2027.527	74.9296			1	411.4776	830.6546		
		162	162				2128	54				411.35	833.85		
	10	197.9846	212.739	87.5092	186.0228	10	2034.459	62.829	2006.305	2065.93	10	972.3696	560.9303	950.2528	413.022
		180	177	76	175		2124.5	62	2075.5	2064.5		981.6	560	952.8	413.05
	30		246.1306			30		1297.337			30		1084.086		
			140					1449					1091.3		
	50			127.0978	109.941	50			2034.181		50			730.0647	
				107	91				2106.5					735.2	
	100	124.3928	156.9248			100	1588.98	58.0668		1174.671	100	644.3332	376.2169		885.305
		105	143				1876	53		1209		644.6	377.7		885.85
batch	pseudoset	10	30	50	100	pseudoset	10	30	50	100	pseudoset	10	30	50	100
	relearn					relearn					relearn				
	1	61.3802	151.1766			1	72.9094	151.1766			1	491.1629	289.2957		
		56	53				53	53				474.7	235.05		
	10	64.867	61.3236	78.0882	79.2608	10	143.184	270.6204	81.5922	701.0954	10	298.6029	581.6812	431.0811	696.3091
		59	54	64	60		53	59	53	312		273.4	570.5	424.3	669.8
	30		413.445			30		79.1648			30		485.2032		
			317					62					484.95		
	50					50			1356.528		50			331.9525	
									1313					320.7	
	100	66.006	85.9946		129.6376	100	361.9262	947.519		53.362	100	359.6483	469.9983		403.4266
		56	76		98		266	79		53		327.25	443.6		376.45

Table 2. POMDP results overall

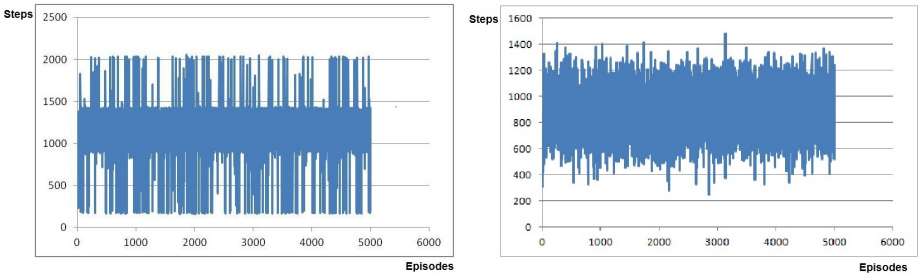

Fig. 3. results of running in POMDP with pseudoset size 100 and relearning gap 100, single run on the left plot and averaged by 10 runs on the right one

The pseudorehearsal approach taken from Frean and Robins decreases this serious context switches and all the agents using this type of pseudorehearsal show nearly the same performance during the all runs, holding around some value with occasional deviations, results much better or worse can be met, but they are rare compared to the results close to this mean.

Stability is maintained for all sets of parameters, while the mean value can differ. Different effect is caused by batch-backpropacation learning using pseudosets: picture of agent's performance is the same as in learning without pseudorehearsal, but efficiency switching occurs only after reinitialization of the pseudoitem vector. While same pseudorehearsal parameters may lead to different agent's behaviour, the efficiency of each set of parameters evaluated by averaging for ten iterations shows that some sets make agent to increase it's performance during the time, while the others do not.

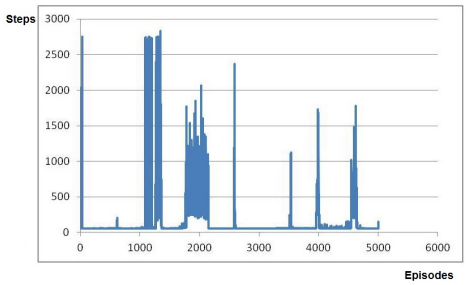

Fig. 4. Example of diverged agent changing policy: batch pseudorehearsal, pseudoset size = 30, no relearning gap - new pseudoset generated after each run

4 Discussion and conclusions

The experiment has shown us that pseudorehearsal can deal with catastrophic interference, but it has its own effects which in some cases cause divergence that worsen performance, so this tool should be used carefully and the parameters - learning rate, pseudoset size and relearning frequency have to be chosen properly to guarantee high performance on the current task.

For the fully observable MDPs pseudorehearsal decreases influence of the catastrophic forgetting if the optimal parameters for the task are known. In the best cases, optimal performance was reached quickly with pseudorehearsal, but the further parameters from the optimal, the worse performance was. In the case if modeling this environment might be too complex - optimal parameters can be defined only empirically before starting learning, which may be unacceptable if the cost of mistake is high.

For partially observable environments all the problems met by fully observable ones remain the same and some additional effects were noted: agent's policy doesn't only converge, diverge or stay random, but also converge to some value with the majority of tries having results in a some range around this value, and strongly deviating runs are more rare and separated by wide gaps of convergence. Another notable effect of the pseudorehearsal in POMDP agents with both pseudorehearsal cases is a significant decrease of the number of steps to converge to the number of steps needed by fully observable agent. If an agent in fully observable environment can converge it converges at about 20-30 runs as with pseudorehearsal, so without, if agent in partially observable environment can converge it converges at 100-250 runs without pseudorehearsal and at 10-30 runs with pseudorehearsal.

Pseudorehearsal is known to be a powerful tool for improving performance of supervised learning agents. We have shown that it can be useful to assist learning even in relatively quickly mixing continuous reinforcement learning tasks, if parameters are chosen correctly. Pseudorehearsal reduces this forgetting effect and maintains stable solutions for longer. While pseudorehearsal may strongly improve agent's performance and accelerate learning, empirical defining of the

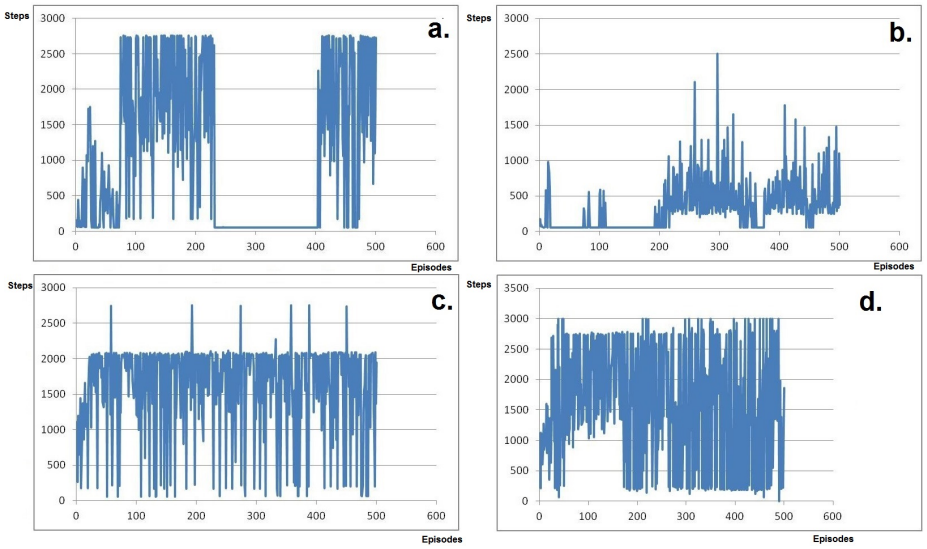

Fig. 5. a., b.500 first stps for two agents without pseudorehearsal; c. one agent with Frean-Robins' pseudorehearsal with pseudoset size = 10 and relearning gap = 100; d. one agent with batch pseudorehearsal with pseudoset size = 100 and relearning gap = 10

optimal pseudoset size and relearning gap is required. One of possible extension of this research would be exploration of mathematical way to figure out this parameters. We will also explore new, more complex reinforcement learning challenges and try more advanced dual network generation of pseudoexamples.

References

1. K. P. Murphy, "A survey of pomdp solution techniques," *environment*, vol. 2, p. X3, 2000.
2. L. Johard and E. Ruffaldi, "A connectionist actor-critic algorithm for faster learning and biological plausibility," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 3903–3909.
3. R. S. Sutton, "Learning to predict by the methods of temporal differences," *Machine learning*, vol. 3, no. 1, pp. 9–44, 1988.
4. J. N. Tsitsiklis and B. Van Roy, "An analysis of temporal-difference learning with function approximation," *IEEE transactions on automatic control*, vol. 42, no. 5, pp. 674–690, 1997.
5. R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, no. 3-4, pp. 229–256, 1992.
6. C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
7. M. Geist and O. Pietquin, "A brief survey of parametric value function approximation," *Rapport interne, Supélec*, 2010.

8. R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press Cambridge, 1998, vol. 1, no. 1.
9. J. Gama, R. Sebastião, and P. P. Rodrigues, "On evaluating stream learning algorithms," *Machine Learning*, vol. 90, no. 3, pp. 317–346, 2013.
10. M. McCloskey and N. J. Cohen, "Catastrophic interference in connectionist networks: The sequential learning problem," *Psychology of learning and motivation*, vol. 24, pp. 109–165, 1989.
11. R. M. French, "Semi-distributed representations and catastrophic forgetting in connectionist networks," *Connection Science*, vol. 4, no. 3-4, pp. 365–377, 1992.
12. R. Coop, A. Mishtal, and I. Arel, "Ensemble learning in fixed expansion layer networks for mitigating catastrophic forgetting," *IEEE transactions on neural networks and learning systems*, vol. 24, no. 10, pp. 1623–1634, 2013.
13. R. Ratcliff, "Connectionist models of recognition memory: constraints imposed by learning and forgetting functions." *Psychological review*, vol. 97, no. 2, p. 285, 1990.
14. G. E. Hinton and D. C. Plaut, "Using fast weights to deblur old memories," in *Proceedings of the ninth annual conference of the Cognitive Science Society*, 1987, pp. 177–186.
15. L.-J. Lin, "Reinforcement learning for robots using neural networks," DTIC Document, Tech. Rep., 1993.
16. S. Adam, L. Busoniu, and R. Babuska, "Experience replay for real-time reinforcement learning control," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 42, no. 2, pp. 201–212, 2012.
17. V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
18. A. Robins, "Catastrophic forgetting, rehearsal and pseudorehearsal," *Connection Science*, vol. 7, no. 2, pp. 123–146, 1995.
19. M. Frean and A. Robins, "Catastrophic forgetting in simple networks: an analysis of the pseudorehearsal solution," *Network: Computation in Neural Systems*, vol. 10, no. 3, pp. 227–236, 1999.
20. B. F. Goodrich, "Neuron clustering for mitigating catastrophic forgetting in supervised and reinforcement learning," Ph.D. dissertation, University of Tennessee, 2015.
21. B. Baddeley, "Reinforcement learning in continuous time and space: Interference and not ill conditioning is the main problem when using distributed function approximators," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 38, no. 4, pp. 950–956, 2008.
22. J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory." *Psychological review*, vol. 102, no. 3, p. 419, 1995.
23. A. Robins and S. McCallum, "The consolidation of learning during sleep: comparing the pseudorehearsal and unlearning accounts," *Neural Networks*, vol. 12, no. 7, pp. 1191–1206, 1999.
24. M. Hattori, "A biologically inspired dual-network memory model for reduction of catastrophic forgetting," *Neurocomputing*, vol. 134, pp. 262–268, 2014.
25. L. Breitwieser, R. Bauer, A. D. Meglio, L. Johard, M. Kaiser, M. Manca, M. Mazzara, F. Rademakers, and M. Talanov, "The biodynamo project: Creating a platform for large-scale reproducible biological simulations," *4th Workshop on Sustainable Software for Science: Practice and Experiences (WSSSPE4)*, 2016.